



US009471674B2

(12) **United States Patent**  
**Ramanathan et al.**

(10) **Patent No.:** **US 9,471,674 B2**  
(45) **Date of Patent:** **Oct. 18, 2016**

(54) **MEDIA FINGERPRINTING AND IDENTIFICATION SYSTEM**

(56) **References Cited**

U.S. PATENT DOCUMENTS

(71) Applicant: **Gracenote, Inc.**, Emeryville, CA (US)

8,195,689 B2 \* 6/2012 Ramanathan ..... G06F 17/30799  
707/747  
8,364,703 B2 \* 1/2013 Ramanathan ..... G06F 17/30799  
707/747

(72) Inventors: **Prashant Ramanathan**, Mountain View, CA (US); **Jose Pio Pereira**, Cupertino, CA (US); **Shashank Merchant**, Sunnyvale, CA (US); **Mihailo M. Stojancic**, San Jose, CA (US)

(Continued)

OTHER PUBLICATIONS

(73) Assignee: **Gracenote, Inc.**, Emeryville, CA (US)

Marco La Cascia and Edoardo Ardizzone, "Jacob: Just a Content-Based Query System for Video Databases", IEEE, 1996, pp. 1216-1219.\*

(\*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

*Primary Examiner* — Cheryl Lewis

(21) Appl. No.: **15/073,858**

(74) *Attorney, Agent, or Firm* — Hultquist, PLLC; Peter H. Priest

(22) Filed: **Mar. 18, 2016**

(65) **Prior Publication Data**

US 2016/0267079 A1 Sep. 15, 2016

(57) **ABSTRACT**

**Related U.S. Application Data**

(63) Continuation of application No. 14/885,110, filed on Oct. 16, 2015, now Pat. No. 9,323,754, which is a continuation of application No. 14/711,054, filed on May 13, 2015, now Pat. No. 9,195,663, which is a

(Continued)


The overall architecture and details of a scalable video fingerprinting and identification system that is robust with respect to many classes of video distortions is described. In this system, a fingerprint for a piece of multimedia content is composed of a number of compact signatures, along with traversal hash signatures and associated metadata. Numerical descriptors are generated for features found in a multimedia clip, signatures are generated from these descriptors, and a reference signature database is constructed from these signatures. Query signatures are also generated for a query multimedia clip. These query signatures are searched against the reference database using a fast similarity search procedure, to produce a candidate list of matching signatures. This candidate list is further analyzed to find the most likely reference matches. Signature correlation is performed between the likely reference matches and the query clip to improve detection accuracy.

(51) **Int. Cl.**  
**G06F 17/30** (2006.01)  
**G06K 9/00** (2006.01)

(52) **U.S. Cl.**  
CPC ..... **G06F 17/30784** (2013.01); **G06K 9/00744** (2013.01)

(58) **Field of Classification Search**  
CPC ..... G06F 17/30097; G06F 17/30112;  
G06F 17/30799; G06K 9/00744  
See application file for complete search history.

**20 Claims, 22 Drawing Sheets**



**Related U.S. Application Data**

continuation of application No. 14/059,688, filed on Oct. 22, 2013, now Pat. No. 9,053,104, which is a continuation of application No. 13/719,603, filed on Dec. 19, 2012, now Pat. No. 8,688,731, which is a continuation of application No. 13/463,137, filed on May 3, 2012, now Pat. No. 8,364,703, which is a continuation of application No. 12/772,566, filed on May 3, 2010, now Pat. No. 8,195,689.

- (60) Provisional application No. 61/185,670, filed on Jun. 10, 2009.

(56)

**References Cited**

U.S. PATENT DOCUMENTS

8,688,731	B2 *	4/2014	Ramanathan .....	G06F 17/30799 707/747
9,053,104	B2 *	6/2015	Ramanathan .....	G06F 17/30799
9,195,663	B2 *	11/2015	Ramanathan .....	G06F 17/30799
9,323,754	B2 *	4/2016	Ramanathan .....	G06F 17/30799
2010/0250510	A1 *	9/2010	Herberger .....	G10H 1/0025 707/705

\* cited by examiner


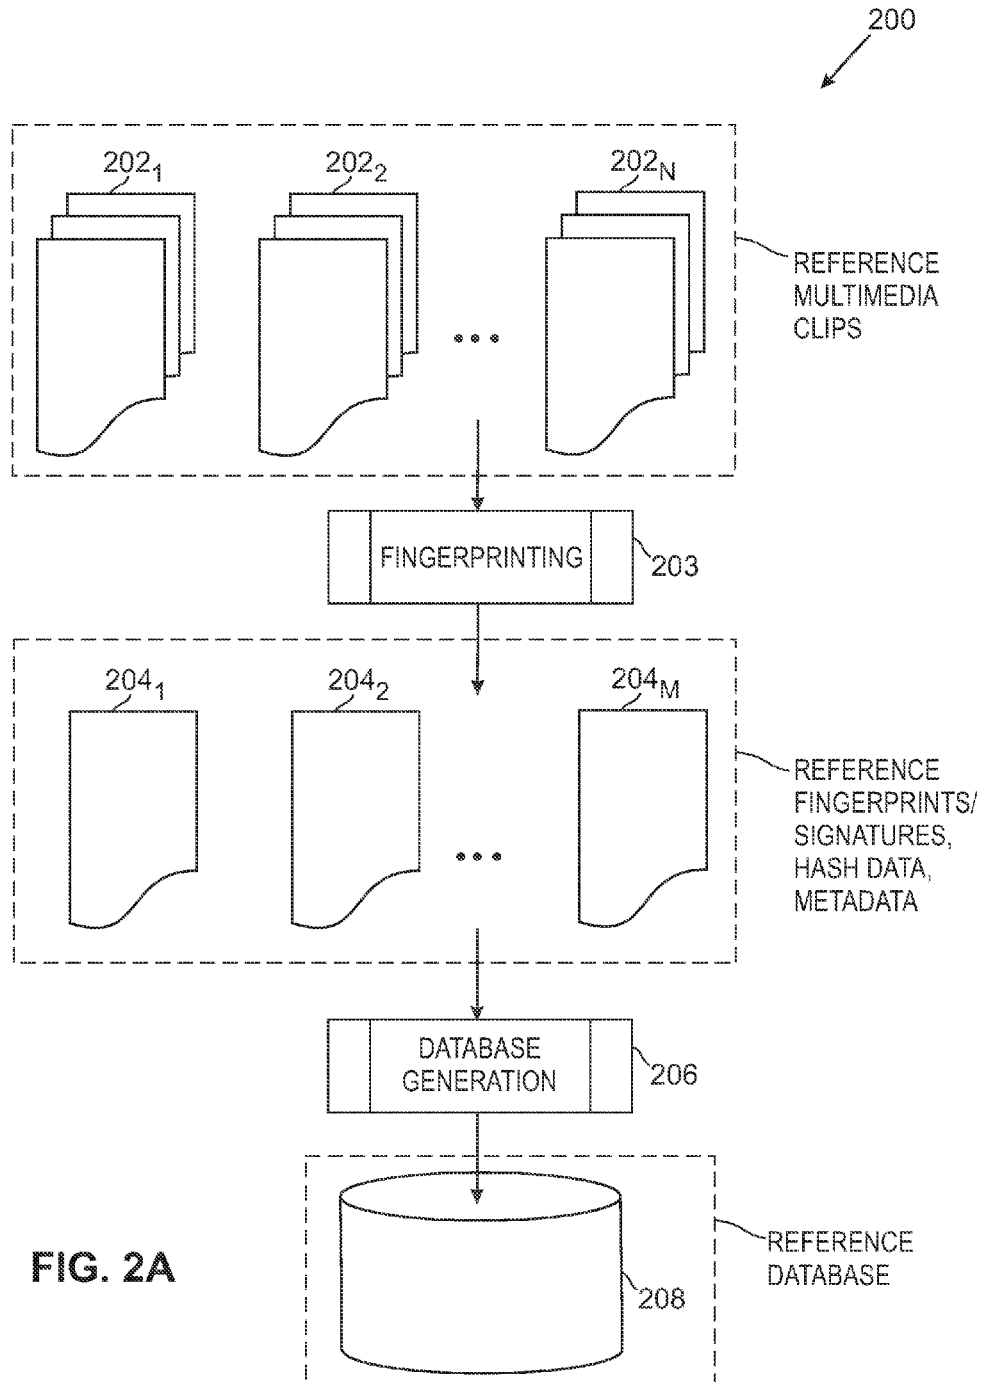


FIG. 1






FIG. 2B




FIG. 2C




FIG. 2D

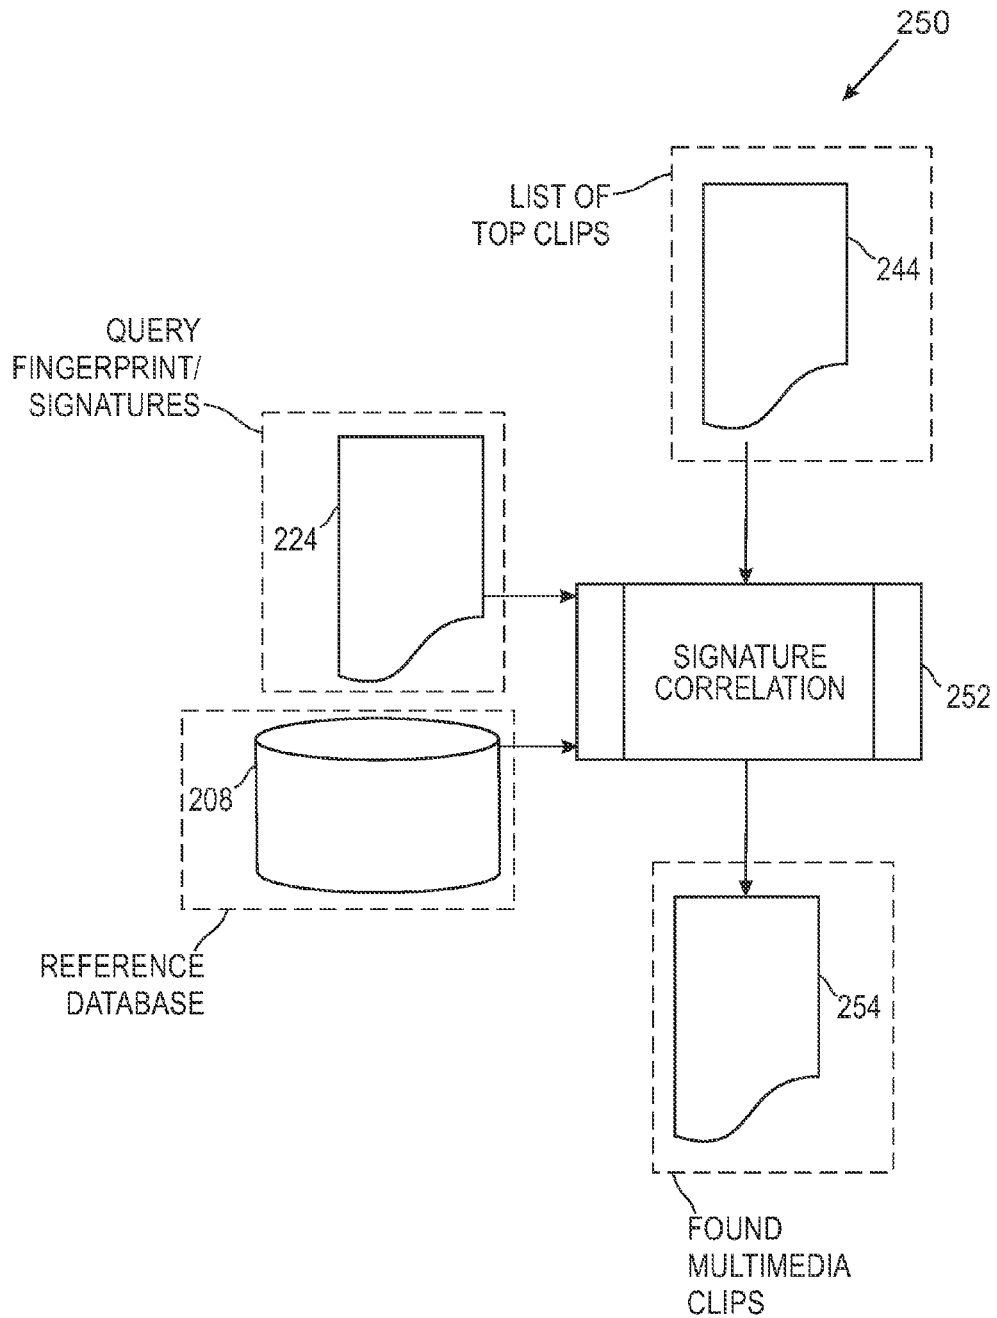



FIG. 2E








FIG. 4




FIG. 5A




FIG. 5B




FIG. 6




FIG. 7




FIG. 8A




FIG. 8B






FIG. 8C




FIG. 8D




FIG. 9A






FIG. 9B






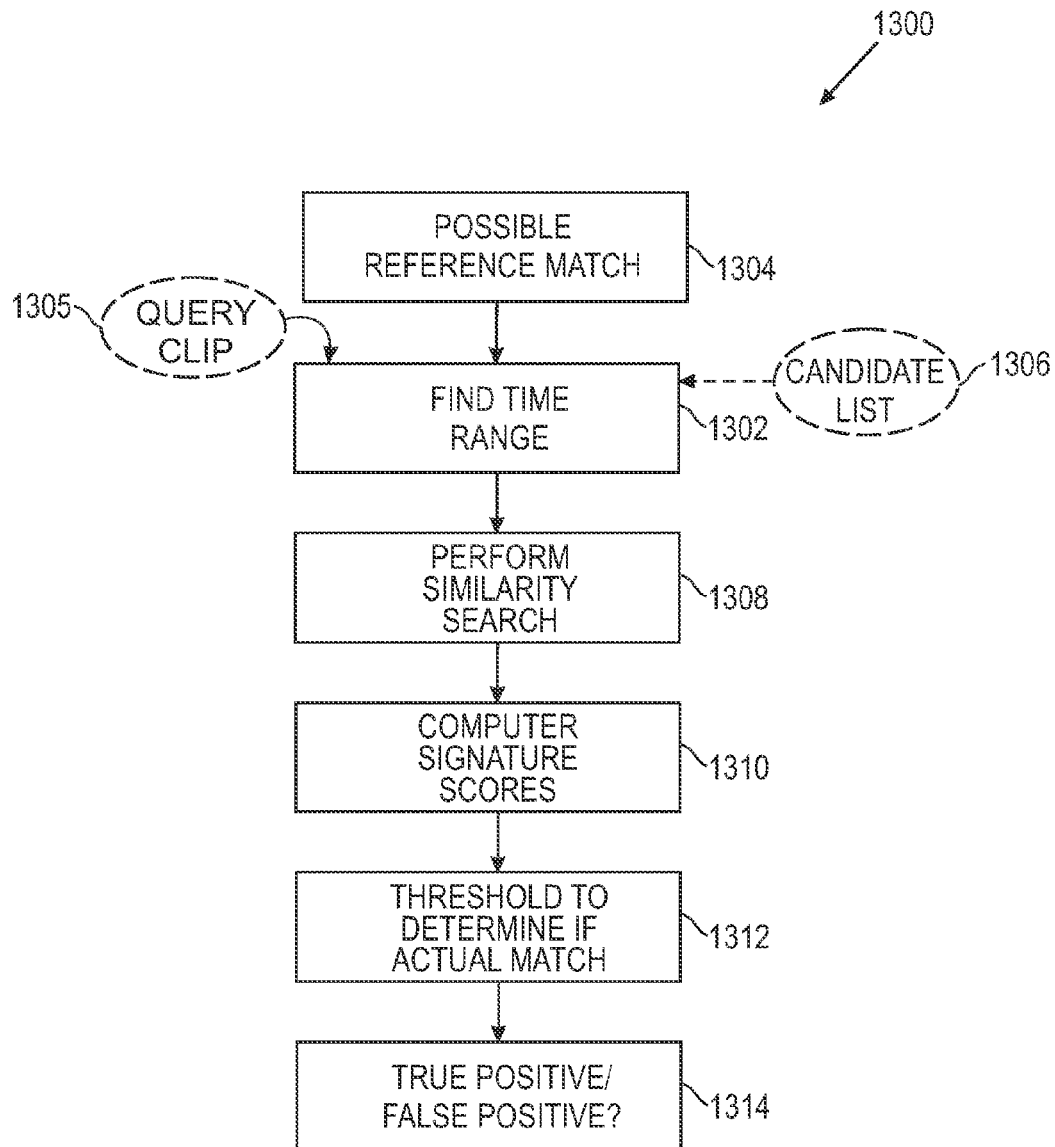


FIG. 12

**FIG. 13**



## MEDIA FINGERPRINTING AND IDENTIFICATION SYSTEM

This application is a continuation of U.S. patent application Ser. No. 14/885,110 entitled "Media Fingerprinting and Identification System" filed Oct. 16, 2015 issued as U.S. Pat. No. 9,323,754, which is a continuation of U.S. patent application Ser. No. 14/711,054 filed on May 13, 2015 issued as U.S. Pat. No. 9,195,663, which is a continuation of U.S. patent application Ser. No. 14/059,688 filed Oct. 22, 2013 issued as U.S. Pat. No. 9,053,104 which is a continuation of U.S. patent application Ser. No. 13/719,603 filed Dec. 19, 2012 issued as U.S. Pat. No. 8,688,731 which is a continuation of U.S. patent application Ser. No. 13/463,137 filed on May 3, 2012 issued as U.S. Pat. No. 8,364,703 which is a continuation of U.S. patent application Ser. No. 12/772,566 filed on May 3, 2010 issued as U.S. Pat. No. 8,195,689 which claims the benefit of U.S. Provisional Patent Application Ser. No. 61/856,670 entitled "A Highly Scalable, Accurate and Distortion-Robust Media Fingerprinting and Identification System" filed on Jun. 10, 2009 all of which are hereby incorporated by reference in their entirety.

### CROSS REFERENCE TO RELATED APPLICATIONS

U.S. application Ser. No. 12/141,163 filed on Jun. 18, 2008 issued as U.S. Pat. No. 8,229,227 entitled "Methods and Apparatus for Providing a Scalable Identification of Digital Video Sequences", U.S. application Ser. No. 12/141,337 filed on Jun. 18, 2008 issued as U.S. Pat. No. 8,171,030 entitled "Method and Apparatus for Multi-dimensional Content Search and Video Identification", U.S. application Ser. No. 12/491,896 filed on Jun. 25, 2009 issued as U.S. Pat. No. 8,385,644 entitled "Digital Video Fingerprinting Based on Resultant Weighted Gradient Orientation Computation", U.S. application Ser. No. 12/612,729 filed Nov. 5, 2009 issued as U.S. Pat. No. 8,189,945 entitled "Digital Video Content Fingerprinting Based on Scale Invariant Interest Region Detection with an Array of Anisotropic Filters", U.S. application Ser. No. 12/788,796 filed on May 24, 2010 issued as U.S. Pat. No. 8,335,786 which claims priority from U.S. Provisional Application Ser. No. 61/181,806 filed on May 28, 2009 entitled "Multi-Media Content Identification Using Multi-Level Content Signature Correlation and Fast Similarity Search", and U.S. application Ser. No. 12/955,416 filed on Nov. 29, 2010 which claims priority from U.S. Provisional Application Ser. No. 61/266,668 filed on Dec. 4, 2009 entitled "Digital Video Content Fingerprinting Using Image Pixel Intensity and Color Information" have the same assignee as the present application, are related applications, and are hereby incorporated by reference in their entirety.

### FIELD OF THE INVENTION

The present invention relates generally to improvements in representing and identifying multimedia information. More particularly, the present invention addresses a scalable and accurate audio and video fingerprinting and identification system that is robust to many classes of audio and video distortions.

### GENERAL BACKGROUND OF THE INVENTION

Media applications which include video and audio database management, database browsing and identification are

undergoing explosive growth and are expected to continue to grow. To address this growth, there is a need for a comprehensive solution related to the problem of creating a video sequence database and identifying, within such a database, a particular video sequence or sequences that are tolerant of media content distortions. Multiple applications include video database mining, copyright content detection for video hosting web-sites, contextual advertising placement, and broadcast monitoring of video programming and advertisements.

Multimedia fingerprinting refers to the ability to generate associated identifying data, referred to as a fingerprint, from the multimedia image, audio and video content. A fingerprint ideally has several properties. First, the fingerprint should be much smaller than the original data. Second, the fingerprint should be designed such that it can be searched for in a large database of fingerprints. Third, the original multimedia content should not be able to be reconstructed from the fingerprint. Fourth, for multimedia content that is a distorted version of another multimedia content, fingerprints of the original and distorted versions should be similar. Examples of some common multimedia distortions include, selecting a clip of video content temporally, cropping the image data, re-encoding the image or audio data to a lower bit-rate, changing a frame rate of the video or audio content, re-recording the multimedia data via some analog medium such as a camcorder in a movie theatre, and changing the aspect ratio of the image content. A fingerprint with the fourth property is deemed to be robust against such distortions.

Such a system of fingerprinting and search is preferable to other methods of content identification. For example, multimedia watermarking changes the multimedia content by inserting watermark data. Unlike multimedia watermarking, fingerprinting does not change the content. Fingerprinting is, however, a very challenging problem.

Increasing demand for such fingerprinting and search solutions, which include standard definition (SD) and high definition (HD) formats of video, requires increasing sophistication, flexibility, and performance in the supporting algorithms and hardware. The sophistication, flexibility, and performance that are desired exceed the capabilities of current generations of software based solutions, in many cases, by an order of magnitude.

### SUMMARY OF THE INVENTION

In one or more of its several aspects, the present invention recognizes and addresses problems such as those described above. The present invention describes a system that is highly scalable to very large multimedia databases, highly accurate in finding a closest matching multimedia clip and not misidentifying a wrong clip, and highly robust to many types of distortion. It is also very efficient in terms of the computer hardware resources necessary to fingerprint and search. To such ends, an embodiment of the invention addresses a method of searching reference multimedia signatures that represent audio and video content of reference multimedia clips for a corresponding query multimedia clip. An N dimension signature based on an N dimension descriptor and a traversal hash signature based on a combination of dimensions of the N dimension descriptor are generated for each feature identified in a set of reference multimedia clips. Each N dimension signature and a link to a corresponding reference multimedia clip are stored at a location in a reference signature database addressable by the traversal hash signature. An N dimension query signature based on an N dimension query descriptor and a traversal hash query

3

signature based on a combination of dimensions of the N dimension query descriptor are generated for a query multimedia clip. The reference signature database is searched using the traversal hash query signature to find similar signatures that are within a specified signature distance, wherein the similar reference multimedia clips are aggregated in a candidate list of closely matching signatures that correspond to similar reference multimedia clips.

Another embodiment of the invention addresses a method of forming accurate descriptors of features in multimedia clips. A scale and a location of the center parameters are determined for each feature of a plurality of features selected from a frame from a multimedia clip. A Bx by By grid is created having Bx\*By blocks located at the center of each feature and each block having a Cx by Cy pixel grid. For each block, functions of multiple intensity differences are computed between each pixel in a block's Cx by Cy pixel grid and another pixel for a specified pixel offset that is greater than 1 and based on the scale and location of center parameters. A first N dimension descriptor is generated as a concatenation of the functions for each block of the Bx\*By blocks. The scale and location of center parameters are individually adjusted to account for an expected error margin due to image distortions to generate at least a second N dimension descriptor as a concatenation of the functions for each block of the Bx\*By blocks as adjusted according to the expected error margin.

Another embodiment of the invention addresses a computer readable medium storing a computer program which causes a computer system to perform a method of searching reference multimedia signatures that represent audio and video content of reference multimedia clips for a corresponding query multimedia clip. An N dimension signature based on an N dimension descriptor and a traversal hash signature based on a combination of dimensions of the N dimension descriptor are generated for each feature identified in a set of reference multimedia clips. Each N dimension signature and a link to a corresponding reference multimedia clip are stored at a location in a reference signature database addressable by the traversal hash signature. An N dimension query signature based on an N dimension query descriptor and a traversal hash query signature based on a combination of dimensions of the N dimension query descriptor are generated for a query multimedia clip. The reference signature database is searched using the traversal hash query signature to find similar signatures that are within a specified signature distance, wherein the similar reference multimedia clips are aggregated in a candidate list of closely matching signatures that correspond to similar reference multimedia clips.

A further embodiment of the invention addresses method for creating a media database. A plurality of signatures are generated from descriptors of selected media content. A signature is selected from the plurality of signatures based on a distribution of bit values in the selected signature and on values of a corresponding descriptor. A traversal hash signature is generated from the corresponding descriptor. The selected signature is stored in a media database at a location within a set of contiguous locations based on the traversal hash signature used as a primary index.

These and other features, aspects, techniques and advantages of the present invention will be apparent to those skilled in the art from the following detailed description, taken together with the accompanying drawings and claims.

#### BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 illustrates a system for media content fingerprinting in accordance with the present invention;

4

FIG. 2A illustrates a reference media database generation process in accordance with the present invention;

FIG. 2B illustrates a query fingerprint generation process in accordance with the present invention;

FIG. 2C illustrates a similarity search process in accordance with the present invention;

FIG. 2D illustrates a candidate video filtering process in accordance with the present invention;

FIG. 2E illustrates a signature correlation process in accordance with the present invention;

FIG. 3 illustrates a multimedia fingerprinting process in accordance with the present invention;

FIG. 4 illustrates a scale-space feature in accordance with the present invention;

FIG. 5A illustrates an exemplary rectangular descriptor box and coordinates for use in generating a global descriptor in accordance with the present invention;

FIG. 5B illustrates an exemplary resampled rectangular descriptor box for use in generating a global descriptor in accordance with an embodiment of the present invention;

FIG. 6 illustrates aspects of a global descriptor gradient computation process in accordance with the present invention;

FIG. 7 illustrates multiple robust descriptors in accordance with the present invention;

FIG. 8A illustrates an exemplary signature generation process in accordance with the present invention;

FIG. 8B illustrates a threshold scale used in a procedure to augment a traversal hash signature in accordance with the present invention;

FIG. 8C illustrates a signature selection procedure which selects signatures for a particular frame of video in accordance with the present invention;

FIG. 8D illustrates a procedure to generate the temporal robustness estimate in accordance with the present invention;

FIG. 9A illustrates an exemplary signature database organized by a primary hash signature index in accordance with the present invention;

FIG. 9B illustrates an exemplary signature database organized by secondary index in accordance with the present invention;

FIG. 10 illustrates an exemplary similarity search process in accordance with the present invention;

FIG. 11 illustrates a query multimedia clip partitioned into smaller segments in accordance with the present invention;

FIG. 12 illustrates an exemplary listing of top matches in accordance with the present invention; and

FIG. 13 illustrates a signature correlation process in accordance with the present invention.

#### DETAILED DESCRIPTION

The present invention will now be described more fully with reference to the accompanying drawings, in which several embodiments of the invention are shown. This invention may, however, be embodied in various forms and should not be construed as being limited to the embodiments set forth herein. Rather, these embodiments are provided so that this disclosure will be thorough and complete, and will fully convey the scope of the invention to those skilled in the art.

It will be appreciated that the present disclosure may be embodied as methods, systems, or computer program products. Accordingly, the present inventive concepts disclosed herein may take the form of a hardware embodiment, a software embodiment or an embodiment combining soft-

5

ware and hardware aspects. Furthermore, the present inventive concepts disclosed herein may take the form of a computer program product on a computer-readable storage medium having computer-usable program code embodied in the medium. Any suitable computer readable medium may be utilized including hard disks, CD-ROMs, optical storage devices, flash memories, or magnetic storage devices.

Computer program code or software programs that are operated upon or for carrying out operations according to the teachings of the invention may be written in a high level programming language such as C, C++, JAVA®, Smalltalk, JavaScript®, Visual Basic®, TSQL, Perl, use of .NET™ Framework, Visual Studio® or in various other programming languages. Software programs may also be written directly in a native assembler language for a target processor. A native assembler program uses instruction mnemonic representations of machine level binary instructions. Program code or computer readable medium as used herein refers to code whose format is understandable by a processor. Software embodiments of the disclosure do not depend upon their implementation with a particular programming language.

The methods described in connection with the embodiments disclosed herein may be embodied directly in hardware, in a software module executed by a processor, or in a combination of the two. A software module may reside in RAM memory, flash memory, ROM memory, EPROM memory, EEPROM memory, registers, hard disk, a removable disk, a CD-ROM, or any other form of storage medium known in the art. A computer-readable storage medium may be coupled to the processor through local connections such that the processor can read information from, and write information to, the storage medium or through network connections such that the processor can download information from or upload information to the storage medium. In the alternative, the storage medium may be integral to the processor.

A fingerprint system should be robust to distortions, while allowing for high matching accuracy and low false positive rate. A fingerprinting system should have fast search speed, especially when searching over a large database of content. The fingerprinting system should be able to return a match for a relatively short clip of query content. The fingerprinting system should also be able to generate fingerprints rapidly when generating a reference database or query fingerprints to be searched. These goals tend to compete with each other. For instance, high search speed and support for short query clips can come at the cost of accuracy or robustness to distortion. Generally, accuracy trade-offs against robustness to distortion.

Due to the large size of such databases and the density of media files, high performance, accurate media identification and search technologies are needed. Furthermore, robust video content identification, management, and copyright protection should be resistant to intentional or unintentional video content change or distortion within the limits of parameters, such as, reasonable view-ability. The overall design should be scalable to efficiently handle very large databases of videos and an arbitrary length of a query sequence.

To provide for such needs, FIG. 1 illustrates a system 100 for both media fingerprinting and identification in accordance with an embodiment of the present invention. The system 100 includes user sites 102 and 103, a server 106, a video database 108, and a video fingerprinting and video

6

as the Internet, a wireless network, or a private network, connects sites 102 and 103 and server 106. Each of the user sites, 102 and 103 and server 106 may include a processor complex having one or more processors, having internal program storage and local user controls such as a monitor, a keyboard, a mouse, a printer, and may include other input or output devices, such as an external file storage device and communication interfaces.

The user site 102 may comprise, for example, a personal computer, a laptop computer, or the like equipped with programs and interfaces to support data input and output and video fingerprinting and search monitoring that may be implemented both automatically and manually. The user site 102, for example, may store programs, such as the video fingerprinting and search system 112 program implementation of a content based video identification process of the present invention or have access to such programs through electronic media, such as may be downloaded over the Internet from an external server, accessed through a universal serial bus (USB) port from flash memory, accessed from disk media of various types, or the like. The system 100 may also suitably include more servers and user sites than shown in FIG. 1. Also, multiple user sites each operating an instantiated copy or version of the video fingerprinting and search system 112 may be connected directly to the server 106 while other user sites may be indirectly connected to it over the network 104.

User sites 102 and 103 may generate user video content which is uploaded over the Internet 104 to a server 106 for storage in a video database 108. The user sites 102 and 103, for example, may also operate a video fingerprinting and video identification system 112 to generate fingerprints and search for video content in the video database 108. The video fingerprinting and video identification system 112 in FIG. 1 is scalable and utilizes highly accurate video fingerprinting and identification technology that checks unknown video content against a database of previously fingerprinted video content, which is considered an accurate or “golden” database. The video fingerprinting and video identification system 112 is different from commonly deployed systems in that it extracts features from the video itself. The video fingerprinting and video identification system 112 allows the server 106 to configure a “golden” database specific to its business requirements. The user site 102 that is configured to connect with the network 104, uses the video fingerprinting and search system 112 to compare video streams against a database of signatures in the video database 108, that has been previously generated.

As used herein, a video fingerprint is a calculated numerical value, alternatively termed a signature, which represents a detected feature of a video frame. In one approach, a feature detection procedure uses a two pass scale space analysis as described in U.S. application Ser. No. 12/612, 729 filed Nov. 5, 2009 entitled “Digital Video Content Fingerprinting Based on Scale Invariant Interest Region Detection with an Array of Anisotropic Filters”, FIGS. 3, 4, 5, 6A, 6B, 7A, 7B, and 8 and pages 16-26 and 34-36. The first pass uses Laplacian-of-Gaussian second order partial derivative bi-level filters to generate interest regions. These interest regions are further refined in the second pass, using anisotropic filters of sampled elliptic-shaped Gaussian second order partial derivative filters. The detected features, found by interpolating local maxima values, are characterized by the following set of parameters ( $s_x$ ,  $s_y$ ,  $x$ ,  $y$ , peak polarity), where  $s_x$ ,  $s_y$  denote the scale in the  $x$  and  $y$  directions,  $(x,y)$  denote the pixel location, and peak polarity denotes whether the feature is a dark or light relative to the

image background. For each feature that is detected in a video frame, one or more descriptors are generated according to a global descriptor procedure described below. For each global descriptor, a pair of signatures, a primary signature and a traversal hash signature, are generated as described in more detail below. A collection of signatures drawn from a single video frame with multiple features generates a compact video frame representation, while a collection of signatures associated with each video frame in a video sequence with multiple frames forms the video sequence representation.

The user site **102** has access to the video database **108**. The video database **108** may store the video archives, as well as all data related to various video content stored in the video database **108**, and a plurality of video fingerprints that have been adapted for use as described herein and in accordance with the present invention. It is noted that depending on the size of an installation, the functions of the video fingerprinting and search system **112** and the capacity and management of the video database **108** may be combined in a single processor system, such as user site **102** or server **106**, running separate program threads for each function.

One embodiment of the present invention addresses the problem of matching a query multimedia clip to a reference multimedia clip from a large collection of reference multimedia clips. The multimedia clip and the reference multimedia clips are processed to generate a compact, searchable set of fingerprints and a multimedia fingerprint database is generated to allow for fast, accurate searching using a query fingerprint to search this multimedia fingerprint database.

At the base of the problem of matching a video clip to a video from a large collection of videos, is the problem of reliably matching two digital images when one of them has undergone some form of distortion, such as geometric and optical distortions. Thus, the problem may be restated as what techniques are required to reliably match a distorted query video frame, chosen from a query video clip, to a corresponding undistorted, original video frame residing in a database of original video material. Since the database of original video material may be very large, many video frames in this database could exhibit substantial statistical similarity. Furthermore, many features within a single video frame may also show statistical similarity in both a visual or optical and algorithmic or geometric sense. This fact makes it difficult to clearly distinguish between many possible feature descriptor matches resulting from a database browsing or search activity when a non-exact, similarity measure is used. Hence, a precise, discriminative, and robust video frame feature characterization is desirable having a feature detector and identifier that are well localized in terms of its location spatially and its spatial extent or scale.

This invention describes a system that is readily scalable to very large multimedia databases, accurate in finding a correct clip and not misidentifying a wrong clip, and robust to many types of distortion. A fingerprinting and search system is used, where the fingerprint for a piece of multimedia content is composed of a number of compact signatures, including traversal hash signatures and associated metadata. The compact signatures and traversal hash signatures are constructed to be easily searchable when scaling to a large database of multimedia fingerprints. The multimedia content is also represented by many signatures that relate to various aspects of the multimedia content that are relatively independent from each other. Such an approach allows the system to be robust to distortion of the multimedia content even when only small portions of the multimedia content are available.

Multimedia, specifically audio and video content, may undergo several different types of distortions. For instance, audio distortions include re-encoding to different sample rates or audio quality. Video distortions include cropping, stretching, re-encoding to a lower quality, and image overlays. While these distortions change the digital representation, the multimedia is perceptually similar to a human listener or viewer. Robustness to these distortions refers to the property that content that is perceptually similar will generate fingerprints that have a small distance according to some distance metric, such as Hamming distance for bit based signatures, and content that is perceptually distinct from one another will generate fingerprints that have a large distance, according to the same distance metric. A search for perceptually similar content, hence, is transformed to a problem of searching for fingerprints that are a small distance away from the desired fingerprints.

The overall architecture and details of a highly scalable and accurate video fingerprinting and identification system is described that is robust to many classes of video distortions. In this system, the fingerprint for a piece of multimedia content is composed of a number of compact signatures, along with traversal hash signatures and associated metadata. The compact signatures and traversal hash signatures are constructed so as to be easily searchable when scaling to a very large database of multimedia fingerprints. The multimedia content is represented by many signatures that relate to various aspects of the multimedia that are relatively independent from each other. This allows the system to be very robust under heavy distortion when only small portions of the multimedia content are available.

The content based multimedia identification system consists of a multimedia fingerprinting process and a search process. In order to identify an unknown multimedia clip, these two components are used as described in more detail below.

FIG. 2A illustrates a reference media database generation process **200** in accordance with the present invention. Reference multimedia clips **202<sub>1</sub>**, **202<sub>2</sub>**, . . . **202<sub>N</sub>** that are relevant to the application at hand are identified. The clips **202<sub>1</sub>**, **202<sub>2</sub>**, . . . **202<sub>N</sub>** refer to distinct pieces of multimedia content. For example, the clips could be from a movie and represent different temporal portions of the movie, including audio and video content, or they could be from different movies. Using a video fingerprinting function **203** of the video fingerprinting and search system **112**, reference signatures **204<sub>1</sub>**, **204<sub>2</sub>**, . . . **204<sub>N</sub>** are generated for the reference multimedia clips **202<sub>1</sub>**, **202<sub>2</sub>**, . . . **202<sub>N</sub>**, respectively, along with hashing data and associated metadata. Different pieces of multimedia content may be fingerprinted independently, leading to a parallelizable system. A clip of multimedia content can be broken down into its temporal components, such as frames in the case of video, or temporal audio segments in case of audio. Even more parallelism is possible by independently fingerprinting at the temporal component level, or frame level.

The set of reference signatures **204<sub>1</sub>**, **204<sub>2</sub>**, . . . **204<sub>N</sub>** created in the video fingerprinting function **203** is organized by database generation function **206** into a reference database **208**. This set of reference signatures is indexed by the generated hashing data, described in further detail below. The associated metadata, also described in further detail below, is stored along with each reference signature. A set of reference signatures may also be indexed in other ways, for instance, by multimedia identifiers. A single multimedia identifier denotes a distinct piece of multimedia content. For

instance, the multimedia clips  $202_1, 202_2, \dots, 202_N$  would each be represented by their own multimedia identifier.

FIG. 2B illustrates a query fingerprint generation process **220** in accordance with the present invention. A user requests identification of an unknown multimedia clip **222**, including audio and video content, also referred to herein as a query multimedia clip **222**. The query multimedia clip **222** is processed by the video fingerprinting function **203** to generate query signatures, hash data, and associated meta-data, known collectively as a query fingerprint **224**, for the unknown multimedia clip **222**.

FIG. 2C illustrates a similarity search process **230** in accordance with the present invention. For each query signature in the query fingerprint **224**, a similarity search function **232** is initiated to find similar signatures in the reference database **208**. The hash data associated with each query signature is used to restrict the similarity search function **232** to a relatively small portion of the reference data, allowing the similarity search to be extremely fast even for large reference databases. Only reference signatures that are “similar” within a distance measure to the query signature are returned. These classified similar reference signatures are added to a candidate list **234**, which contains identifying information regarding which reference multimedia clip  $202_1, 202_2, \dots, 202_N$  the similar reference signature belongs.

FIG. 2D illustrates a candidate video filtering process **240** in accordance with the present invention. The video filtering process **240** analyzes the candidate list **234** for the most likely matches in the reference database **208**. The candidate list **234** is sorted in top multimedia clips function **242** to find the top most likely matching multimedia clips. The resulting data is stored in a list of top clips **244**. The list of top clips **244** includes a multimedia identifier for the similar reference multimedia clip. The matching multimedia clip might be only for a specific portion of the multimedia clip, for instance, a particular time segment, such as seconds **93** to **107** of a video sequence, or spatial locations, such as top left quadrant of the video frames. The temporal identification and spatial locations are also included in the list of top clips.

FIG. 2E illustrates a signature correlation process **250** in accordance with the present invention. The list of top clips **244** is selected for correlation. For each of the clips in the top clips list **244**, a set of signatures is accessed from the reference database **208** given the multimedia identifier and any temporal and/or spatial information. A query could correspond to “all signatures for video number ABC from time 10.4 seconds to 25.7 seconds in the bottom-right quadrant of the frame.” These signatures are not restricted to have matching traversal hash signatures. The set of query signatures is now compared against this set of reference signatures, using the signature correlation procedure **252**. For each query signature, a score is derived based on the distances to the closest signatures in the reference database signature subset, and the number of matches. This could include, for instance, the average distance. These scores are combined, for example, the scores are averaged, for the entire set of query signatures to give an overall score for the database clip. Based on whether this score is over a threshold, the database clip is determined to be a true positive or a false positive. The signature correlation procedure **252** is repeated for all clips in the top clips list **244**, to produce a list of matching reference videos, since there may be more than one, if similar content is repeated in the database. FIG. 2A fingerprinting function **203** and FIG. 2B belong to the

multimedia fingerprinting system, while FIG. 2A database generation function **206** and FIGS. 2C-2E belong to the search system.

FIG. 3 illustrates a multimedia fingerprinting process **300** in accordance with the present invention. A video clip is decoded from a multimedia source, such as an motion picture experts group (MPEG) compressed audio and video stream, into individual frames and a subset of the frames are selected for processing. For each selected frame, features are detected using a scale-space feature detector. An N-dimensional descriptor is generated based on the image pixels corresponding to the scale and position values from the feature descriptor. An M-dimensional signature and K-dimensional traversal hash signature is generated from descriptor.

The multimedia fingerprinting process **300** encompasses the fingerprinting function **203** shown in FIG. 2A and the query fingerprint generation process **220** of FIG. 2B. As shown in FIG. 3, the multimedia fingerprinting process **300** is comprised of a decoding step **304**, a selecting frame step **306**, a detecting audio and video features step **308**, a describing the detected audio and video features step **310**, a generating signatures and hash data step **312**, and a selecting signatures step **314**.

The decoding step **304** comprises decoding a multimedia stream that may be accessed in compressed form from a distributed source to produce raw or unprocessed image and audio data. This data, which is associated with an identified playout timestamp, can be combined with other contiguous data to form a “frame” of data. For example, for image data, a frame could be a single image, or group of images, decoded from the multimedia stream. For audio data, a frame could be several milliseconds worth of decoded audio.

The selecting frame step **306** comprises selecting a subset of all available frames for further content analysis. Image and audio data exhibit a great deal of temporal coherence which makes it possible to select a limited number of frames from the entire set of data and obtain similar accuracy as if all frames were used. The search system may also be more efficient when fewer signatures are used. The selecting frame step **306** uses an algorithm to select the subset of all available frames to represent the entire multimedia sequence, as described in more detail below.

The detecting audio and video features step **308** comprises further analyzing each selected frame to find audio and video features that should be reproducible under an expected suite of distortions. A reproducible feature implies a high likelihood of having corresponding features in the original content and the distorted content. For an image, this may involve doing scale-space analysis or using a blob detector, to find an x and y location and an x and y scale of the identified blobs in the image.

The describing the detected audio and video features step **310** comprises generating a compact descriptor for each of the features detected in the previous step. For image data, the compact descriptor, such as a vector of real values, could be based on intensity and gradient values in a spatial area surrounding a detected blob. Generally, such a compact descriptor has on the order of 64 or 128 components, based on spatial subdivision and various image spatial and temporal derivatives of zeroth, first and higher orders.

The generating signatures and hash data step **312** comprises, for each generated descriptor of a particular feature of a frame of multimedia, generating an associated signature and a signature hash. The signatures generated typically quantize a descriptor’s real value and generate a 1-bit, 2-bit, 4-bit, or 8-bit representation for each of the descriptor’s

11

components. For a descriptor that is 64-dimensional, we can use a single threshold value to generate a 64-bit signature. By summing multiple descriptor dimensions, for instance, 4 dimensions, and thresholding the resulting vector, a 16-bit signature hash value can be generated to associate with this signature.

The selecting signatures step **314** comprises selecting signatures that may have greater information content about the multimedia content. For instance, bit-signatures that have nearly all 0's or all 1's tend not to be very unique to a particular multimedia content. Hence, signatures with a count of 0's or 1's that fall below some threshold may be removed. Signatures may also be rejected based on descriptor information, for the descriptor from which the signature is derived. By summing the descriptor values and rejecting those with a sum below a threshold, signatures are eliminated that have corresponding descriptors with their energy concentrated in only a few dimensions.

The multimedia fingerprinting process **300** is described in more detail beginning with the frame selection step **306**. The frame selection step **306** includes both static and dynamic selection of frames. In static frame selection, frames are selected in a specified frame selection period of N frames, where, for example, every  $N^{th}$  frame is selected without particular regard to frame content. This approach is used for certain applications where processing of every frame, would be difficult due to computational load. In dynamic frame selection, frames are selected based on their audio and video content. Other aspects and further details regarding frame selection may be found in U.S. application Ser. No. 12/141,163 filed on Jun. 18, 2008 entitled "Methods and Apparatus for Providing a Scalable Identification of Digital Video Sequences", FIGS. 3 and 12 and pages 14-16, 25, and 26, and U.S. application Ser. No. 12/491,896 filed on Jun. 25, 2009 entitled "Digital Video Fingerprinting Based on Resultant Weighted Gradient Orientation Computation", FIG. 8 and pages 20-22. For dynamic frame selection, a frame selection period is also specified. The frame selection algorithm and period affects the total processing time, the number of signatures generated, as well as the accuracy of the results.

The detecting audio and video features step **308** involves finding some aspect of the multimedia content for the selected frame that should be reproducible under an expected suite of distortions. In accordance with an embodiment of the present invention, an advantageous approach for finding reproducible image features is described. It is noted that this approach can also be adapted to finding reproducible audio features. Both image and audio features may also be described by a compact descriptor.

Image features, such as blobs in an image, may be identified by location and scale in an x and y direction by using a form of scale-space analysis, also referred to herein as blob detection. Conceptually, scale-space analysis applies Gaussian smoothing filters at a particular scale to the original image, and computes a determinant of the Hessian, to produce a response image for that scale. By appropriately selecting peaks from these response images, x and y locations and scales may be inferred.

Further details regarding scale-space analysis may be found in U.S. application Ser. No. 12/612,729 filed Nov. 5, 2009 entitled "Digital Video Content Fingerprinting Based on Scale Invariant Interest Region Detection with an Array of Anisotropic Filters", FIGS. 3, 4, 5, 6A, 6B, 7A, 7B, and 8 and pages 16-26 and 34-36. Other feature detection techniques may also be alternatively used in the fingerprinting system of the present invention.

12

The multimedia fingerprinting process **300** continues with step **310** to generate one or more feature descriptors. A rectangular region of interest in the image is computed based on the scale and location of the detected feature. "Global gradient" values, a more general form of gradient values, are computed in the x and y directions over the pixels of the region of interest. The descriptor is computed according to a nonlinear weighting of the global gradient values, producing a result which is accumulated into +x, -x, +y, -y bins, for multiple regions in the region of interest. The end result is an M-dimensional descriptor.

An example of a detected scale space feature is shown in FIG. 4. FIG. 4 illustrates a scale-space feature **400** in accordance with the present invention. In FIG. 4, an image feature is represented by dashed outline **409**, which could be a contiguous area of near-constant colored pixel that is, for instance, a human face, or part of a picture on a wall. For this image feature, a corresponding image blob **402** is located in an image space **404** by (centerx, centery) coordinates **406** and with a particular  $s_x$  **407** and  $s_y$  **408** scale reference. A scale-space feature has a few others attributes that describe it relative to the surrounding image pixels. For example, a peak polarity value is a binary value indicating whether a feature is a dark blob on a light background, or a light blob on a dark background. A peak strength value represents a measure of how much a blob stands out from its background. For example, a light gray blob on a dark gray background has a lower peak strength value, as compared to a black blob on a white background which has a higher peak strength value.

The step **310** for describing the detected audio and video features includes generating a global descriptor. FIG. 5A illustrates an exemplary rectangular descriptor box and coordinates **500** for use in generating a global descriptor **500** in accordance with the present invention. Other aspects and further details regarding descriptors, such as a phase descriptor, may be found in U.S. application Ser. No. 12/612,729 filed Nov. 5, 2009 entitled "Digital Video Content Fingerprinting Based on Scale Invariant Interest Region Detection with an Array of Anisotropic Filters", FIGS. 9, 10, and 11 and pages 26-28, 38, and 39.

The detecting audio and video features step **308** generates four output values which include (centerx, centery) coordinates **406**,  $s_x$  **407**, and  $s_y$  **408** which represent the center of an image blob **504** and a spatial extent of the blob. A rectangular box **506** is placed on the image blob with center at the blob (centerx, centery) coordinates **406** and spatial extent ( $Ns_x$  **502**,  $Ms_y$  **503**), where N and M are multiplication factors that determines the size of the neighborhood around the blob that the descriptor considers. This rectangular box **506** is subdivided into a Bx by By grid, with, for example, Bx=4 and By=4 providing a 4x4 grid to create Bx\*By=16 blocks

FIG. 5B illustrates an exemplary resampled rectangular descriptor box **515** for use in generating a global descriptor in accordance with an embodiment of the present invention. In FIG. 5B, an exemplary rectangular descriptor box **522** corresponding to **506** in FIG. 5A, is centered at an interest point **530**, with Bx\*By=16 sub-regions, where Bx and By both equal 4. In more detail and with specified feature parameters (centerx, centery,  $s_x$ ,  $s_y$ ) a rectangular descriptor box **522** is placed on the image with center at (x, y) interest point **530** and with spatial extent of  $Ns_x$ ,  $Ms_y$ , where N and M are multiplication factors which determine the size of the neighborhood around the interest point under consideration. The rectangular descriptor box **522** is resampled to generate a resampled grid **524** with (Bx\*By)(Rx\*Ry), where Rx and

Ry equal 5 in this example, pixel regions, such as region 528 for block 526. For each block, such as the block 526, a Rx\*Ry pixel grid 521 with associated Rx\*Ry re-sampled pixel values in region 528 is computed as follows. First, lines that define the Rx\*Ry sub-grid are determined, where these lines are given with sub-pixel accuracy and are not necessarily aligned with the pixel boundaries 521. In the x direction, the lines evenly subdivide the block 521 horizontally and are named Lnx0, Lnx1, . . . , Lnx(Rx). Likewise, in the y direction, the lines evenly subdivide the block 521 vertically and are named Lny0, Lny1, . . . , Lny(Ry). The locations of the pixels contained within rectangle 521 are shown as dots. The pixels contained within each of the Rx\*Ry sub-grid areas are combined to produce a resampled pixel. This is done by summing the intensity values and dividing by the number of pixels, within each of the Rx\*Ry sub-grid areas, yielding Rx\*Ry pixel intensity values. Note that there may be different number of pixels for some or all of the Rx\*Ry sub-grid areas. There may also be no pixels in a Rx\*Ry sub-grid area, in which case the resampled intensity value is taken as zero.

This Rx\*Ry pixel sub-region, computed for each of the Bx\*By blocks of the rectangular box drawn around an interest point, is subsequently used to generate a descriptor and, based on this descriptor, a signature for that interest point.

FIG. 6 illustrates aspects of a global descriptor gradient computation process 600 in accordance with the present invention. Each Rx\*Ry pixel grid that is computed for each of the Bx\*By blocks is used to generate 4 values. These values are “global gradient” values computed according to the following procedure. For each pixel in an Rx\*Ry grid, the differences from a pixel to several other pixels are computed. FIG. 6 illustrates the pixels on which these differences are taken, for example, from the reference pixel 602, to the target pixels 603, 604, 605 and 606. These target pixels may be chosen arbitrarily. In this example, they have been chosen to be to the right 2 pixels and 4 pixels, as well as down 2 pixels and 4 pixels. Other pixel positions are also possible. Based on these pixels, four difference values are computed, including difference A 607 computed between pixel 603 and pixel 602, difference B 608 computed between pixel 604 and pixel 602, difference C 609 computed between pixel 605 and pixel 602, difference D 610 computed between pixel 606 and pixel 602. Pixels, such as pixels 604 and 606, might not exist in a particular Rx\*Ry pixel grid for pixels near to the right or bottom borders, so such pixels will access the appropriate pixel in a neighboring Rx\*Ry pixels grid. If a bordering Rx\*Ry pixel grid does not exist, an additional block or blocks might need to be computed along the right and bottom borders of the rectangular box 506. This means, for instance, that instead of a Bx by By grid computed, a (Bx+1) by (By+1) grid is computed, with reference pixels within the original Bx by By grid and target pixels potentially accessing the larger (Bx+1) by (By+1) grid.

These difference values are combined into one or more values, using a weighted average of the values. For example, an “x-gradient” is computed as the weighted sum of A and B,

$$xgrad = w_1 * A + w_2 * B,$$

and, similarly, a “y-gradient” is computed as the weighted sum of C and D,

$$ygrad = w_3 * C + w_4 * D.$$

The weights,  $w_1$  to  $w_4$ , may depend on many factors, but in one embodiment, they depend on an agreement of signs of

the difference values A and B, and C and D. For instance, if A and B are of the same sign, then  $w_1=1.1$  and  $w_2=0.2$ , whereas if A and B are of different values, then  $w_1=1.2$  and  $w_2=0.0$ . Different weights can be used for the x and y directions. The weights can also depend on decision factors other than the agreement of signs.

For example, an xgrad value and a ygrad value may be computed for each of the pixels associated with the Rx\*Ry pixel grid 521. From this set of values, four sums are computed as global gradient values:

1.  $-1 * \text{sum of xgrad for each pixel, where xgrad is negative}$
2.  $\text{sum of xgrad for each pixel, where xgrad is positive}$
3.  $-1 * \text{sum of ygrad for each pixel, where ygrad is negative}$
4.  $\text{sum of ygrad for each pixel, where ygrad is positive}$

In total, a descriptor for the rectangular box 506 consists of 4\*Bx\*By values that is a concatenation of the 4 “global gradient” values for each of the Bx\*By blocks covered by the descriptor. By design, these values are all positive. For example, for global gradient values 1 and 3 in the above list, the sums are of negative values, to produce a negative sum, which is made positive by multiplying by  $-1$  and for global gradient values 2 and 4, the sums are of positive values, producing a positive sum.

Some image distortions may cause the feature detection algorithm to have errors in the detected scale values or location. If the error is large enough, this might result in differences in the descriptor, and subsequently in the signature, that are too large for matching to work correctly. To improve the chances of matching, that is, make the fingerprint more robust with respect to certain image distortions, the following embodiment of generating multiple different descriptors based on the same detected feature may be used. These multiple descriptors may be generated, for example, by perturbing the scale value and/or the location value. For example, the scalex multiplication factor, such as N of FIG. 5A, or the scaley multiplication factor, such as M of FIG. 5A, or both are adjusted to account for an expected error margin of  $\pm D\%$ , wherein D % represents effects of image distortions.

FIG. 7 shows an image feature 714, which could be a contiguous area of near-constant colored pixel that is, for instance, a human face, or part of a picture on a wall. This is detected as a blob at location (cx,cy) 710 and scale (sx,sy) represented by ellipse 702. These location and scale values are used in the global descriptor generation procedure previously described above. FIG. 7 shows that two additional descriptors may be generated by first multiplying the scale by 0.9 and 1.1 to generate the representative ellipses 706 and 704, respectively. For these two descriptors, the location is kept constant and the scale is perturbed by the 0.9 and 1.1 multipliers. Alternatively, the scale may be kept constant and the location is moved by a delta amount in the x and y directions (dx,dy) to move the representative ellipse 702 to a new center location (cx+dx,cy+dy) 712 with representative ellipse 708. An additional descriptor for the representative ellipse 708 may be generated using this new center location.

In addition, these multiple descriptors can be generated either for reference clips, the reference side, or for query clips, the query side, or both. By increasing the number of descriptors on either the reference side or the query side the chances of detecting a match are improved, as long as the multiple descriptors cover the errors in the feature detection. Using multiple descriptors on both the reference side and query side tends to further improve the chances for matching. While some multimedia content may experience an

15

increase in the number of false matches with use of multiple descriptors, generally the use of multiple descriptors improves the chances of detecting matches. Using multiple descriptors on the reference side increases the size of the reference database, while using multiple descriptors on the query side increases the query fingerprint size. While the use of multiple descriptors on both the reference side and the query side may lower the search speed, detecting matches is generally improved. In one approach, these increases may be compensated through use of parallel processing techniques.

The technique of generating multiple descriptors is applicable to other descriptor generation procedures, in addition to the global descriptor procedure described above. A descriptor generation procedure that takes as input a scale-space feature with a specified location and scale, generating a compact N dimension vector for a rectangular area of pixels, may also be used.

To enable efficient search capabilities, the descriptors are processed to generate a compact representation that is referred to as a signature. These signatures are implemented as bit vectors. In one embodiment, each dimension of an M-dimensional descriptor generates a single bit in an output signature. For example, a 64-bit signature would be generated for a 64-dimensional descriptor. In other embodiments, two or more bits may be used per dimension of the descriptor. A traversal hash signature is generated to facilitate a fast lookup mechanism. In one embodiment, multiple dimensions of the M-dimensional descriptor are combined to generate a single bit in the traversal hash signature. In other embodiments, multiple dimensions of the M-dimensional descriptor are combined to generate multiple bits in the traversal hash signature.

The signature bits are produced by quantizing each dimension of the descriptor independently. First, the descriptor is normalized in some way, such as made into a unit norm. For example, a unit norm vector may be generated by dividing each component of the vector by the square root of the sum of the squares of each component, using a technique commonly referred to as the L2-norm. For a one-bit-per-dimension signature, the value of each dimension of the normalized descriptor is compared against a pre-specified threshold value. The resulting bits in the signature represent whether the descriptor dimension is above (1) or below (0) the threshold. In another embodiment, which may be more efficient, the descriptor is not normalized, but rather the threshold is scaled according to an L2-norm of the descriptor. In yet another embodiment, the threshold may be simply set to an average value of all the components of the descriptor.

FIG. 8A illustrates an exemplary signature generation process **800** in accordance with the present invention. All positive global descriptor values **802** are shown as vertical bars with a threshold **804** shown as a dashed line. The threshold **804** is either set according to prior empirical data or determined by the descriptor values. Descriptor values greater than the threshold **804** produce a '1' and less than or equal to the threshold **804** produce a '0' generating signature **806**.

Choosing the threshold **804** may be done empirically. The goal is to generate unique signatures, which are more likely unique when the numbers of ones and zeros are approximately equal. A threshold that is set too high may generate signatures that are mostly zeros, whereas a threshold that is set too low may generate signatures that are mostly ones. The threshold should be chosen so as to avoid either extreme.

16

In order to facilitate fast search, a mechanism to search over only a small part of the database is required. We enable this by generating a shorter signature, which acts as a direct hash value for the primary signature. This shorter signature is termed a "traversal hash signature" and is associated with a set of signatures that tend to be similar to one another. The traversal hash signature may also be generated at the same time as the primary signature. For example, for a 64-dimensional descriptor, in addition to a 64-bit signature, such as signature **806**, an H-bit traversal hash signature **808** may be generated, where  $H < 64$ . This traversal hash signature **808** may be based on combining the values of various dimensions into a single combined value and then, similar to the signature generation process, thresholding the combined values with a traversal hash threshold **810**. For example with  $H=16$ , the values of the dimensions in groups of four are combined by either taking the sum of the four values, the average of the four values, or the sum of the square of the four values. For the global descriptor values **802**, a meaningful group of four is the four values for a block which when combined produce the combined descriptor values **812**. For rectangular blocks, such as the rectangular block **506**, there are 16 blocks with each block corresponding to one bit in the traversal hash signature **808**.

A traversal hash signature may be augmented with additional bits that relate to an underlying feature, or perhaps to another description of the image data. For an underlying feature, additional bits may be added that identify the spatial location of the underlying feature, the scale of the underlying feature, or the peak polarity of the underlying feature. For example, the spatial location of the underlying feature may be described by two bits representing the quadrant where the center of the underlying feature is located. A quadrant may be indicated by having one bit indicate top or bottom half and the other bit indicate left or right, such that the quadrant is identified with the two bits. The scale of the feature may be quantized to either one or two bits, by setting one or more thresholds.

FIG. 8B illustrates a threshold scale **838** used in a procedure to augment a traversal hash signature in accordance with the present invention. From the scale values  $s_x$  and  $s_y$ , a normalized scale value  $s$  may be computed according to the equation:  $s = \sqrt{s_x^2 + s_y^2}$ , represented on the threshold scale **838** as a value on a linear scale **820**. This is evaluated against multiple thresholds. If  $s$  is between thresholds  $T_0$  **822** and  $T_1$  **824**, with threshold  $T_0$  **822** inclusive, then the scale hash bits are set to "00" **830**; if  $s$  is between thresholds  $T_1$  **824** and  $T_2$  **826**, with threshold  $T_1$  **824** inclusive, then the scale hash bits are set to "01" **832**; if  $s$  is between thresholds  $T_2$  **826** and  $T_3$  **828**, with threshold  $T_2$  **826** inclusive, then the scale hash bits are set to "10" **834**; if  $s$  is greater than or equal to threshold  $T_3$  **828**, then the scale hash bits are set to "11" **836**. The peak polarity of the underlying feature may be indicated by the feature detection algorithm which indicates whether a feature based on a contrast polarity threshold is a light coloring on a dark background or vice-versa using one bit to represent this characteristic.

The additional bits to augment the traversal hash signature may be based on other descriptors of image data or other data, such as texture or motion. For example, an intensity histogram may be derived for the pixel area **506** in FIG. 5A corresponding to the feature. The intensity histogram may be heavily quantized to produce a small number of bits. Another approach would be to use motion analysis to indicate whether a blob is moving from frame to frame. One



17

bit can indicate whether there is motion or not. The approximate direction of motion may also be quantized to a small number of bits.

The overall hash traversal hash signature is a concatenation of the traversal hash signature generated by appropriately combining and thresholding the descriptor, as illustrated in FIG. 8A, with additional bits generated from scale, as illustrated in FIG. 8B, x/y, peak polarity, or the like.

An advantageous feature of the present invention is the ability to select signatures for inclusion in the fingerprint database and eliminate certain signatures that do not improve identification accuracy and lead to false positives. Various criteria, such as the distribution of ones and zeros in the primary signature and the traversal hash signature, are used to determine whether a signature should be included in the database. Good signature selection may reduce the size of the fingerprint database, reduce the false positive rate and improve or maintain the overall accuracy. In a similar manner, the number of signatures produced for a query multimedia clip may be reduced. As described in more detail below, these advantageous features may be achieved by examining both the signatures and descriptors, individually and in aggregate. This examination is based on multiple criteria, such as descriptor information, signature information, temporal robustness, and spatial diversity.

FIG. 8C illustrates a signature selection procedure 860 which selects signatures for a particular frame of video in accordance with the present invention. The inputs for a frame of video are a set of signatures and associated traversal hash signatures 866, a set of descriptors 864 associated with the set of signatures 866 for the frame, a set of scale, scalex, centerx, centery values 862 associated with the set of signatures 866; and sets of signatures and associated traversal hash signatures 868 for a previous frame. For each signature in the set 866, there is a corresponding descriptor in the set 864 from which the signature is derived, as well as a corresponding set of scale and location values in the set 862, from which the descriptor is derived. The sets of signatures 868 do not necessarily have any correspondence to the signatures in the set 866. Various computations described below are performed based on the inputs above.

For each descriptor in the set 864, descriptor information 872 is computed. The descriptor information 872 is associated with the corresponding signature. Descriptor information may include a sum of the descriptor values after normalization to unit norm. Descriptors that tend to produce false positives have energy concentrated in only a few dimensions and hence have a lower "descriptor information" value. As an example, this tends to happen, for instance, where a descriptor pixel region of interest contains a bright blob on a dark background of constant intensity. Such regions of interest can be confused with other regions with similar characteristics and lead to false matches. These regions of interest lead to descriptors with energy concentrated in a few dimensions, with those dimensions corresponding to the blob center.

For each primary signature in the set 866, signature information 874 is computed. The signature information 874 is associated with the corresponding signature. Signature information may include a count of the number of 1's in a signature. Signatures that tend to produce false positives have only a few bits turned on and hence have lower "signature information" value. This tends to happen for the same reasons as the descriptor. While "signature information" and "descriptor information" are related, their thresholds can be picked independently to provide more flexibility.

18

For each signature in the set 866, a temporal robustness value 876 is computed, using the sets of signatures 868. The computed temporal robustness value 876 is associated with the corresponding signature in the set 866. Temporal robustness may include a measure of whether a feature appears in a following frame. The frame lag between a frame under consideration and the following frame may be one frame, the immediately following frame, or more than one frame. An estimate is generated to indicate whether a feature appears in a following frame based solely on the signature and traversal hash signature. FIG. 8D illustrates a procedure 881 to generate the temporal robustness estimate in accordance with the present invention. A set of signatures and traversal hash signatures for the current frame 882, is compared against all signatures and traversal hash signatures from following frame 883 by computing the bit distance between the respective signatures in traversal hash signature comparison computation 884 and the bit distance between the respective signatures in signature comparison computation 885. The bit distances are computed for each signature in 882, and the bit distances are compared against a robustness estimate threshold computation 886 to produce a set of temporal robustness values 887, one value corresponding to each signature in 882. Specifically, a signature is considered "temporally robust" if there exists at least one pair of signature and traversal hash signature found in the traversal hash signature comparison computation 884 and in the signature comparison computation 885 with a signature bit distance  $\leq R$  and a traversal signature bit distance  $\leq Q$ . The values R and Q are integer-valued parameters that can be specified. In one embodiment, a temporal robustness value is either 0 or 1, based on whether the signature is "temporally robust".

For each feature scale and location value in the set 862 in FIG. 8C, a spatial diversity value 870 is computed by comparing its own value against the entire list of feature scale and location values in the set 862. The computed spatial diversity value 870 is associated with the corresponding signature. Spatial diversity may include a measure of whether a feature is located close to other features. For example, the distance between any pair of feature locations can be computed for the set 862, and used to compute the average distance between a feature center location and the M near features' center locations, where M may be pre-specified. This average distance, or "closeness" indicates the density of features in terms of location. Ideally, the features would be spread out and those features that are concentrated too close together would be penalized. By incorporating this "closeness" criterion, features that are spread out are preferred, according to a procedure described in more detail below.

Each signature in the set 866 is now associated with the results of the computations in blocks 870, 872, 874, and 876. The results are input to the signature sorting and selection block 878. Signatures that fail to meet some minimum level of criteria are eliminated. For descriptor information, a minimum threshold may be used. Likewise, for signature information, another minimum threshold may be used. For temporal robustness, Q and R bit distance parameters may be set, and signatures that are not considered temporally robust according to this criterion are eliminated. Similarly, for spatial diversity, a threshold can be established for "closeness", and signatures that are too close based on this threshold are eliminated.

Signatures may also be sorted, for example, according to the descriptor information, the signature information and the spatial diversity, descending, in that order. The sorting

provides a list of signatures in approximate order of strength. Then, all but the T top signatures are eliminated, where T is a parameter that can be pre-specified. The result is a set of signatures **880** that satisfy the minimum criteria and/or are the top T signature according to some criteria.

The selected signatures including the signature and traversal hash signature pairs are stored along with associated metadata such as multimedia id, playout timestamp, and scale, and x and y location values to create a signature record that is inserted into the reference signature database. The signature database consists of the set of all signature records for all corresponding reference multimedia along with two indices that are used for different types of access of the signature records. These indices are a hash signature index and an index based on the multimedia id and playout timestamp. There may be a single set of signature records pointed to by both indices, or two sets of signature records, with identical data, for example, one for each index.

FIG. 9A illustrates an exemplary signature database **900** organized by a primary hash signature index in accordance with the present invention. The signature records for all the multimedia content that is to be put into the database is collected together and grouped by the traversal hash signature. At this stage of processing, the number of signatures that belong to particular traversal hash signature is known so the memory for the signature records can be allocated and signature records may be stored in memory. It is advantageous for search operations that the signature records for a particular traversal hash signature be stored contiguously. The set of signatures belonging to a traversal hash signature is called a cluster. For example, 100 signatures and corresponding traversal hash signatures may be generated having 50 signatures with a traversal hash signature A, 30 signatures with a traversal hash signature B and 20 signatures with a traversal hash signature C. Thus, the 100 signatures are organized into three clusters, cluster A with 50 signatures, cluster B with 30 signatures and cluster C with 20 signatures, that are stored in memory.

For each traversal hash signature, the number of signatures and a pointer to the location where the corresponding signature records begin is stored for processing. Since the space of traversal hash signatures may be relatively small, for example, a 16-bit traversal hash signature implies a maximum 65,536 entries, all possible traversal hash signatures can be indexed in an array. A traversal hash signature array **902** is shown in FIG. 9A storing one element for each possible traversal hash signature. The index into the traversal hash signature array **902** is the integer interpretation of the traversal hash signature as a binary number. Thus, given a traversal hash signature, direct addressing into the array retrieves the number of matching signatures and where corresponding signature records are located, such as a link reference to a list of signature records **904** or **906**. In FIG. 9A, for example, traversal hash signature "0101001010101010" is located at entry **908**, which links to the array of signature records **904**, and traversal hash signature "1011010100101010" is located at entry **910**, which links to the array of signature records **906**. Each entry in traversal hash signature array **902**, such as entries **908** and **910**, have an additional field included in the entry that stores the link reference address (LRA) to a signature record array, such as signature record arrays **904** and **906**.

FIG. 9B illustrates an exemplary signature database **950** organized by secondary index in accordance with the present invention. The secondary index supports accesses of the signature records corresponding to a multimedia id, from a starting playout time to an ending playout time. The sec-

ondary index is useful for the signature correlation process, discussed in more detail below. The secondary index is based on a hash table **952** organized by multimedia id. For example, a hash table entry **958** corresponding to multimedia id "vo102340910" points to a linked list **960**. Each entry in the linked list, for example, **962** and **964**, points to an array of signature records, for example, **954** and **956**. The signature records are stored in playout timestamp order within each array, and the linked list is organized with arrays in playout timestamp order. Alternatively, instead of **954** and **956** representing an array of signature records, they may represent an array of pointers to signature records, already stored according to the primary index. Hence, the secondary index does not store the signature records again, but rather points to the signature records of the first index. This is the case where both indices shared a common set of signature records, as mentioned above. Such an alternative approach can be more space efficient, but may be slower to access due to signatures that are randomly distributed relative to the second index.

The secondary index allows for searches such as 'all reference signatures for multimedia id "vo102340910" between time period 817000 milliseconds and 1048500 milliseconds'. In this case, the entry **958** corresponding to "vo102340910" is first accessed, retrieving the linked list **960**, starting with the linked list entry **962**. The array of signatures **954** can be examined as to whether the time period 817000 to 1048500 milliseconds is contained. Any matching signatures are returned. The subsequent entries on the linked list, for instance **964**, are processed, until there are no more entries, to ensure that all matching signatures within the time period are retrieved. This search can be made faster by keeping the starting and end timestamp of an array of signatures within the corresponding linked list entry itself, so that the signatures themselves do not need to be accessed unnecessarily. Also, the requested signature record for the desired starting timestamp is found by binary search through the array. From there, signature records in the array are processed linearly until the desired ending timestamp is reached.

The video database structure may be stored either in the local computer's main memory (its RAM) or on a hard disk drive. One embodiment is to store the signature database in main memory as access speeds are significantly faster. The trade-off is the smaller capacity of main memory versus the hard drive.

A query multimedia clip is processed to generate signatures as described above and using the process of similarity search against a reference database, the most likely candidate clips may be identified. The beginning step in a similarity search is to find all near-matching signature records in the video database for a query multimedia clip. The similarity search process begins with generating signature and traversal hash signature pairs for the query multimedia clip. A playout timestamp is also associated with each signature and traversal hash signature pair. The traversal hash signature is used to dramatically speed up this portion of the search. The result of this initial search is a list of matching signatures and their associated metadata called the candidate list. In a following step, the number of times a multimedia clip appears in the candidate list is counted to generate a frequency list of multimedia clip identifications (ids), which is then sorted from most frequent to least frequent. The most frequent clip ids on this list are the likely matching multimedia clips. In a correlation step, the top matching multimedia clips are analyzed to find the closest matching clip. In the correlation step, signatures of the top

21

matching clips are used instead of matching traversal hash signatures, and a score is computed based on the amount of signature bit error. This score is compared against prespecified thresholds to decide whether the corresponding multimedia clip is a good match to the query clip. Thus, one or more closest matching multimedia clips from the reference database may be chosen for presentation as a result of the search.

FIG. 10 illustrates an exemplary similarity search process 1000 in accordance with the present invention. For each signature and traversal hash signature pair 1002 in a list of query signatures 1004, a similarity search is performed given the parameters of signature distance R and traversal hash signature distance Q. A traversal hash signature distance of Q=0 means that only an exact match of the traversal hash signature is allowed, thus only a single cluster needs to be searched, which is one of the embodiments of the present invention. A traversal hash signature distance of Q=1 means that traversal hash signatures which are a maximum of 1 bit distance from the specified traversal hash signature are searched. For an K-bit hash signature, this would mean that a total of K+1 clusters need to be searched including the original cluster, as well as the K clusters that correspond to exactly one of the bits flipped and all other bits unchanged. The number of clusters needed to be searched increases exponentially with Q.

In the case of Q=0, the traversal hash signature "0101010100101010" 1003 in the signature and traversal hash signature pair 1002 obtained from the list of query signatures 1004 is used to look up entry 1012 in the traversal hash signature array 1006 and obtain a set of corresponding signature records 1008. This set contains the potentially matching signatures with the same traversal hash signature. The query signature 1009, shown as 31 bits of a 64 bit query signature "1011101001001010100010101010101 . . ." in the pair 1002 is now compared with each signature in the set 1008. Specifically, the distance between the query signature and each signature in the set to be compared, is computed. This computation may be made more efficient by using bit operands and a bit-count lookup table. The Hamming or bit distance between two 64-bit numbers can be computed by performing an XOR on the two numbers and then count the number of 1's in the result. The counting step may be implemented with a lookup table that maps an 8-bit or 16-bit number to the number of 1's contained. The number of one's in the result of the XOR is then the sum of the lookup table results for all of the 8-bit or 16-bit chunks contained in that XOR result. If the distance computed is less than or equal to the specified signature distance R, then this signature record is inserted into a candidate list 1010.

An entry in the candidate list 1010 consists of a query clip timestamp, a matching database clip multimedia id, a matching database clip timestamp, and a signature bit distance. For each signature and traversal hash signature pair, several entries may be added to the candidate list 1010. Similarity search is repeated for each signature and traversal hash signature pair 1002 in the list of query signature records 1004 corresponding to the query multimedia clip. Depending on the length of the query clip, the candidate list 1010 may have hundreds or thousands of entries.

FIG. 11 illustrates a query multimedia clip partitioned into smaller segments 1100 in accordance with the present invention. The query multimedia clip 1102 may additionally be broken into smaller overlapping segments 1104-1106, and similarity search and a top matches and correlation procedure performed on the smaller segments. This partitioning allows more time granularity in the multimedia clip identi-

22

fication, as well as greater accuracy in identifying matching multimedia clips. In one embodiment, the entire query multimedia clip is broken into small clip segments that overlap by as much as 50% with the neighboring clip segments.

FIG. 12 illustrates an exemplary listing of top matches 1200 in accordance with the present invention. The candidate list 1202 that corresponds to a query multimedia clip, or some small segment of it, may be processed to produce the most likely matching multimedia clips. In a first step, the number of times a database multimedia clip appears in the candidate list is counted, to generate a frequency list of multimedia clip id's and the frequency of that clip's appearance in the candidate list. In a second step, the frequency list is then sorted in descending order of frequency to generate a sorted frequency list 1204 representing the top matches.

For application scenarios that require low computational burden, for instance where a search needs to be extremely fast, or where the query clips are known to have little or no image distortion, due to the nature of the application, the correlation step can be omitted. In this case, the most frequent database multimedia clip, which would be represented by the first multimedia clip id in the sorted frequency list 1204, may be simply output as the identity of the unknown query clip. Alternatively, for other scenarios, by taking the top V most frequent multimedia clips, the first V multimedia clip id's in the sorted frequency list 1204, the candidate list has been focused in on the most likely matches. Next, signature correlation processing is used to eliminate potential false matches and ensure, with greater accuracy, that a correct matching multimedia clip has been found.

FIG. 13 illustrates a signature correlation process 1300 in accordance with the present invention. Further information regarding a procedure for using signature correlation to eliminate false positives and get time-accurate matches may be found in U.S. Provisional Application Ser. No. 61/181,806 filed on May 28, 2009 entitled "Multi-Media Content Identification Using Multi-Level Content Signature Correlation and Fast Similarity Search", FIGS. 1A, 2A, 2B, 2C, 3A, 3B, 4-6, pages 11-26. The signal correlation process 1300, in the context of a signature database, is described next. One of the goals of the signature correlation process 1300 is to verify that the signatures of a potentially matching reference multimedia clip 1304 do indeed correspond with the signatures of the query multimedia clip 1305.

In a first step 1302, a time range of the potentially matching reference multimedia clip is determined from the reference database. For example, if the potentially matching reference multimedia clip in the reference database is subdivided into smaller clips, then an approximate time range may be determined. Otherwise, and potentially in addition, the matching time range can be inferred from the candidate list 1306 by analyzing a correspondence in time between the database clip and the query clip, using a Hough transform.

In a second step 1308, with a matching time range and the associated database clip, corresponding signatures are accessed from the signature database using the secondary index as illustrated in FIG. 9B. Generally, the corresponding signatures are a small set of signatures. Since only the matching time range is accessed, the size of the set of signatures remains small, even as the size of the database grows, allowing the correlation step to scale with the database size. The signatures for the matching time range of the reference are now compared with the query signatures. For each query signature, a similarity search is performed, with a specified signature distance, within this small set of

23

signatures from the reference database. This step is used to determine the bit distance from each query signature to each reference signature, and is used in the next step to determine scores.

In a third step **1310**, for each query signature, a score is derived based on the distances to the closest signatures in the reference database signature subset, and the number of matches. For instance, the score assigned to a query signature may be the minimum bit distance to all the signatures in the reference signature set, with a particular maximum allowable value. The score for each query signature may be averaged for the entire set of query signatures to give an overall score for the particular reference clip match. According to the steps above, a low score implies low bit distance between the query and reference, and hence a good match.

In a fourth step **1312**, the overall score is subjected to a threshold test which determines whether the database clip is considered a match **1314**. This threshold must be appropriately selected, based on application requirements, to trade-off between a high match rate and a low false positive rate.

Those of skill in the art will appreciate that based on the present disclosure additional alternative systems and methods for scalable, accurate and distortion robust media fingerprinting and identification of digital video sequences may be determined in accordance with the disclosed principles of the present invention. Thus, while particular embodiments and applications of the present invention have been illustrated and described, it is to be understood that the invention is not limited to the precise construction and components disclosed herein and that various modifications, changes and variations which will be apparent to those of ordinary skill in the art may be made in the arrangement, operation and details of the method and apparatus of the present invention disclosed herein without departing from the spirit and scope of the invention as defined in the appended claims. For example, in another embodiment, an alternate method of feature detection other than the two-pass scale space approach, is used. A feature detection method that provides the center location and spatial extent of the feature, which may be described as a "scale space" method, is feasible. In yet another embodiment, an alternate method for generating a descriptor is used other than the "global descriptor" method presently described. The "phase descriptor" method is one such alternative, and it is described in U.S. application Ser. No. 12/612,729 filed Nov. 5, 2009 entitled "Digital Video Content Fingerprinting Based on Scale Invariant Interest Region Detection with an Array of Anisotropic Filters", FIGS. 9, 10, and 11 and pages 26-28, 38, and 39. Other approaches that generate a compact N dimension vector for a rectangular area of pixels are applicable.

We claim:

1. A method for generating a digital representation of an image feature in multimedia content, the method comprising:

receiving from a source of multimedia content a sequence of image data that is combined into frames, time stamped, and stored as received frames on a user personal device;

reducing the received frames to selected frames that represent the multimedia content to be processed on the user personal device based on temporal coherence of the multimedia content across multiple time stamped frames;

determining center coordinates (centerx, centery) of an interest point in an image feature and a scale in an

24

X-direction (Sx) and a scale in a Y-direction (Sy) for the image feature detected in a frame of the selected frames;

establishing a grid of blocks located at the (centerx, centery) coordinates of the image feature with spatial extent of JSx by KSy, wherein each block contains a grid of pixels and wherein J and K are multiplication factors which determine size of the neighborhood around the interest point;

computing M functions of pixel intensity values of the pixels in each grid of pixels to produce M function results for each block; and

generating a first digital representation of the image feature as a concatenation of the M function results for each block of the established grid of blocks, wherein the concatenation produces the first digital representation as a first N dimension descriptor of the image feature.

2. The method of claim 1 further comprising:

adjusting individually the Sx, Sy, and (centerx, centery) coordinates by a pre-specified error percentage to account for image distortions and to generate for each adjustment at least a second N dimension descriptor as a concatenation of the M functions for each block of the established grid of blocks as adjusted according to the pre-specified error percentage.

3. The method of claim 2, wherein the first N dimension descriptor and the second N dimension descriptor are stored in a media content database on the user personal device as part of a set of N dimension descriptors that represent other image features in the selected frames that represent the multimedia content.

4. The method of claim 2, wherein the first N dimension descriptor and the second N dimension descriptor are used to search a media content database on the user personal device having sets of N dimension descriptors for different multimedia content.

5. The method of claim 1 further comprising:

computing the M functions as multiple intensity differences between each pixel in a block's grid of pixels and another pixel at a pre-specified pixel offset.

6. The method of claim 1 further comprising:

computing, for each pixel in the block's grid of pixels, weighted gradient values in the x direction and weighted gradient values in the y direction that are combined to produce the M function results.

7. The method of claim 1 further comprising:

applying a first threshold to the first N dimension descriptor to produce a first N dimension signature; combining dimensions of the first N dimension descriptor to produce a first combined dimension descriptor, and applying a second threshold to the first combined dimension descriptor to produce a first traversal hash signature.

8. The method of claim 7 further comprising:

storing the first N dimension signature and the second N dimension signature at a location in a media content database on the user personal device addressable by the first traversal hash signature.

9. The method of claim 7 further comprising:

searching a media content database on the user personal device using the first traversal hash signature in a query to find reference N dimension signatures that are within a specified signature distance from the first N dimension signature that accompanies the first traversal hash signature in the query.

25

10. A method for generating a digital representation of an image feature in multimedia content, the method comprising:

- receiving from a source of multimedia content a sequence of image data that is combined into frames, time stamped, and stored as received frames on a user personal device;
- reducing the received frames to selected frames that represent the multimedia content to be processed on the user personal device based on temporal coherence of the multimedia content across multiple time stamped frames;
- determining center coordinates (centerx, centery) of an image feature and a scale in an X-direction (Sx) and a scale in a Y-direction (Sy) for the image feature detected in a frame of the selected frames;
- establishing a Bx by By grid of pixels located at the (centerx, centery) coordinates of the image feature with spatial extent of JSx by KSy, wherein the Bx by By grid of pixels is sub-divided into blocks, each block having a subset of pixels from the Bx by By grid of pixels and wherein J and K are multiplication factors which determine size of a neighborhood around the (centerx, centery) coordinates;
- generating a first digital representation of the image feature as a concatenation of results from M gradient functions calculated for each block of the established Bx by By grid of pixels, wherein the concatenation produces the first digital representation as a first N dimension descriptor of the image feature; and
- applying a first threshold to the first N dimension descriptor to generate a binary first N dimension signature representation of the image feature.

11. The method of claim 10 further comprising:
- adjusting individually the Sx, Sy, and (centerx, centery) coordinates by a pre-specified error percentage to account for image distortions and to generate for each adjustment at least a second N dimension descriptor as a concatenation of results from the M gradient functions for each block of the established Bx by By grid of pixels as adjusted according to the expected error percentage; and
  - applying the first threshold to the second N dimension descriptor to generate a binary second N dimension signature representation of the image feature.

12. The method of claim 11, wherein the binary first N dimension signature and the binary second N dimension signature are stored in a media content database on the user personal device as part of a set of binary N dimension signatures that represent other image features in the selected frames that represent the multimedia content.

13. The method of claim 11, wherein the binary first N dimension signature and the binary second N dimension signature are stored in a media content database on a server remote from the user personal device as part of a set of binary N dimension signatures that represent other image features in the selected frames that represent the multimedia content.

14. The method of claim 11, wherein the binary first N dimension signature and the binary second N dimension signature are used to search a media content database on the user personal device having sets of binary N dimension signatures for different multimedia content.

15. The method of claim wherein 11, the binary first N dimension signature and the binary second N dimension signature are used to search a media content database on the

26

user personal device having sets of binary N dimension signatures for different multimedia content; and

upon determining the search of the media content database on the user personal device was not successful, searching for the binary first N dimension signature and the binary second N dimension signature on a media content database on a remote server.

16. The method of claim 10 further comprising:

- computing multiple intensity differences for each pixel in each block having a subset of pixels from the Bx by By grid of pixels to another pixel in both x and y directions for a pre-specified pixel offset; and
- computing weighted sum combinations of the multiple intensity differences as x-gradient values and as a y-gradient values to produce the results from the M gradient functions.

17. The method of claim 10 further comprising:

- combining dimensions of the first N dimension descriptor to produce a first combined dimension descriptor; and
- applying a second threshold to the first combined dimension descriptor to produce a first traversal hash signature.

18. The method of claim 17 further comprising:

- storing the first N dimension signature and the second N dimension signature at a location in a media content database on the user personal device addressable by the first traversal hash signature.

19. The method of claim 17 further comprising:

- searching a media content database on the user personal device using the first traversal hash signature in a query to find reference N dimension signatures that are within a specified signature distance from the first N dimension signature that accompanies the first traversal hash signature in the query.

20. A computer readable non-transitory medium storing a computer program which causes a computer system to perform a method for generating a digital representation of an image feature in multimedia content, the method comprising:

- receiving from a source of multimedia content a sequence of image data that is combined into frames, time stamped, and stored as received frames on a user personal device;

- reducing the received frames to selected frames that represent the multimedia content to be processed on the user personal device based on temporal coherence of the multimedia content across multiple time stamped frames;

- determining center coordinates (centerx, centery) of an image feature and a scale in an X-direction (Sx) and a scale in a Y-direction (Sy) for the image feature detected in a frame of the selected frames;

- establishing a Bx by By grid of pixels located at the (centerx, centery) coordinates of the image feature with spatial extent of JSx by KSy, wherein the Bx by By grid of pixels is sub-divided into blocks, each block having a subset of pixels from the Bx by By grid of pixels and wherein J and K are multiplication factors which determine size of a neighborhood around the (centerx, centery) coordinates;

- generating a first digital representation of the image feature as a concatenation of results from M gradient functions calculated for each block of the established Bx by By grid of pixels, wherein the concatenation produces the first digital representation as a first N dimension descriptor of the image feature; and

**27**

applying a first threshold to the first N dimension descriptor to generate a binary first N dimension signature representation of the image feature.

\* \* \* \* \*

**28**